

Faster Information Gathering in Ad-Hoc Radio Tree Networks

Marek Chrobak*

Kevin Costello[†]

December 8, 2015

Abstract

We study information gathering in ad-hoc radio networks. Initially, each node of the network has a piece of information called a *rumor*, and the overall objective is to gather all these rumors in the designated target node. The ad-hoc property refers to the fact that the topology of the network is unknown when the computation starts. Aggregation of rumors is not allowed, which means that each node may transmit at most one rumor in one step.

We focus on networks with tree topologies, that is we assume that the network is a tree with all edges directed towards the root, but, being ad-hoc, its actual topology is not known. We provide two deterministic algorithms for this problem. For the model that does not assume any collision detection nor acknowledgement mechanisms, we give an $O(n \log \log n)$ -time algorithm, improving the previous upper bound of $O(n \log n)$. We also show that this running time can be further reduced to $O(n)$ if the model allows for acknowledgements of successful transmissions.

1 Introduction

We study the problem of *information gathering* in ad-hoc radio networks. Initially, each node of the network has a piece of information called a *rumor*, and the objective is to gather all these rumors, as quickly as possible, in the designated target node. The nodes communicate by sending messages via radio transmissions. At any time step, several nodes in the network may transmit. When a node transmits a message, this message is sent immediately to all nodes within its range. When two nodes send their messages to the same node at the same time, a *collision* occurs and neither message is received. Aggregation of rumors is not allowed, which means that each node may transmit at most one rumor in one step.

The network can be naturally modeled by a directed graph, where an edge (u, v) indicates that v is in the range of u . The ad-hoc property refers to the fact that the actual topology of the network is unknown when the computation starts. We assume that nodes are labeled by integers $0, 1, \dots, n-1$. An information gathering protocol determines a sequence of transmissions of a node, based on its label and on the previously received messages.

Our results. In this paper, we focus on ad-hoc networks with tree topologies, that is the underlying ad-hoc network is assumed to be a tree with all edges directed towards the root, although the actual topology of this tree is unknown.

*Department of Computer Science, University of California at Riverside, USA. Research supported by NSF grants CCF-1217314 and CCF-1536026.

[†]Department of Mathematics, University of California at Riverside, USA. Research supported by NSA grant H98230-13-1-0228

We consider two variants of the problem. In the first one, we do not assume any collision detection or acknowledgment mechanisms, so none of the nodes (in particular neither the sender nor the intended recipient) are notified about a collision after it occurred. In this model, we give a deterministic algorithm that completes information gathering in time $O(n \log \log n)$. Our result significantly improves the previous upper bound of $O(n \log n)$ from [5]. To our knowledge, no lower bound for this problem is known, apart from the trivial bound of $\Omega(n)$ (since each rumor must be received by the root in a different time step).

In the second part of the paper, we also consider a variant where acknowledgments of successful transmissions are provided to the sender. All the remaining nodes, though, including the intended recipient, cannot distinguish between collisions and absence of transmissions. Under this assumption, we show that the running time can be improved to $O(n)$, which is again optimal for trivial reasons, up to the implicit constant.

While we assume that all nodes are labelled $0, 1, \dots, n-1$ (where n is the number of vertices), our algorithms' asymptotic running times remain the same if the labels are chosen from a larger range $0, 1, \dots, N-1$, as long as $N = O(n)$.

Related work. The problem of information gathering for trees was introduced in [5], where the model without any collision detection was studied. In addition to the $O(n \log n)$ -time algorithm without aggregation – that we improve in this paper – [5] develops an $O(n)$ -time algorithm for the model with aggregation, where a message can include any number of rumors. Another model studied in [5], called *fire-and-forward*, requires that a node cannot store any rumors; a rumor received by a node has to be either discarded or immediately forwarded. For fire-and-forward protocols, a tight bound of $\Theta(n^{1.5})$ is given in [5].

The information gathering problem is closely related to two other information dissemination primitives that have been well studied in the literature on ad-hoc radio networks: broadcasting and gossiping. All the work discussed below is for ad-hoc radio networks modeled by arbitrary directed graphs, and without any collision detection capability.

In *broadcasting*, a single rumor from a specified source node has to be delivered to all other nodes in the network. The naïve ROUNDROBIN algorithm (see the next section) completes broadcasting in time $O(n^2)$. Following a sequence of papers [6, 18, 2, 3, 21, 11] where this naïve bound was gradually improved, it is now known that broadcasting can be solved in time $O(n \log D \log \log(D\Delta/n))$ [10], where D is the diameter of G and Δ is its maximum in-degree. This nearly matches the lower bound of $\Omega(n \log D)$ from [9]. Randomized algorithms for broadcasting have also been well studied [1, 19, 11].

The *gossiping* problem is an extension of broadcasting, where each node starts with its own rumor, and all rumors need to be delivered to all nodes in the network. The time complexity of deterministic algorithms for gossiping is a major open problem in the theory of ad-hoc radio networks. Obviously, the lower bound of $\Omega(n \log D)$ for broadcasting [9] applies to gossiping as well, but no better lower bound is known. It is also not known whether gossiping can be solved in time $O(n \text{polylog}(n))$ with a deterministic algorithm, even if message aggregation is allowed. The best currently known upper bound is $O(n^{4/3} \log^4 n)$ [16] (see [6, 25] for some earlier work). The case when no aggregation is allowed (or with limited aggregation) was studied in [4]. Randomized algorithms for gossiping have also been well studied [11, 20, 7]. Interested readers can find more information about gossiping in the survey paper [15].

Connections to other problems. This research, as well as the earlier work in [5], was motivated by the connections between information gathering in trees and other problems in distributed

computing involving shared channels, including gossiping in radio networks and MAC contention resolution.

For arbitrary graphs, assuming aggregation, one can solve the gossiping problem by running an algorithm for information gathering and then broadcasting all rumors (as one message) to all nodes in the network. Thus an $O(n \text{ polylog}(n))$ -time algorithm for information gathering would resolve in positive the earlier-discussed open question about the complexity of gossiping. Due to this connection, developing an $O(n \text{ polylog}(n))$ -time algorithm for information gathering on arbitrary graphs is likely to be very difficult – if possible at all. We hope that developing efficient algorithms for trees, or for some other natural special cases, will ultimately lead to some insights helpful in resolving the complexity of the gossiping problem in arbitrary graphs.

Some algorithms for ad-hoc radio networks (see [4, 17], for example) involve constructing a spanning subtree of the network and disseminating information along this subtree. Better algorithms for information gathering on trees may thus be useful in addressing problems for arbitrary graphs.

The problem of *contention resolution* for multiple-access channels (MAC) has been widely studied in the literature. (See, for example, [12, 22, 14] and the references therein.) There are in fact myriad of variants of this problem, depending on the characteristics of the communication model. Generally, the instance of the MAC contention resolution problem involves a collection of transmitters connected to a shared channel (e.g. ethernet). Some of these transmitters need to send their messages across the channel, and the objective is to design a distributed protocol that will allow them to do that. The information gathering problem for trees is in essence an extension of MAC contention resolution to multi-level hierarchies of channels, where transmitters have unique identifiers, and the structure of this hierarchy is not known.

2 Preliminaries

We now provide a formal definition of our model and introduce notation, terminology, and some basic properties used throughout the paper.

Radio networks with tree topology. In the paper we focus exclusively on radio networks with tree topologies. Such a network will be represented by a tree \mathcal{T} with root r and with $n = |\mathcal{T}|$ nodes. The edges in \mathcal{T} are directed towards the root, representing the direction of information flow: a node can send messages to its parent, but not to its children. We assume that each node $v \in \mathcal{T}$ is assigned a unique label from $[n] = \{0, 1, \dots, n-1\}$, and we denote this label by $\text{label}(v)$.

For a node v , by $\deg(v)$ we denote the *degree* of v , which is the number of v 's children. For any subtree X of \mathcal{T} and a node $v \in X$, we denote by X_v the subtree of X rooted at v that consists of all descendants of v in X .

For any integer $\gamma = 1, 2, \dots, n-1$ and any node v of \mathcal{T} define the γ -height of v as follows. If v is a leaf then the γ -height of v is 0. If v is an internal node then let g be the maximum γ -height of a child of v . If v has fewer than γ children of γ -height equal g then the γ -height of v is g . Otherwise, the γ -height of v is $g+1$. The γ -height of v will be denoted by $\text{height}_\gamma(v)$. In case when more than one tree are under consideration, to resolve potential ambiguity we will write $\text{height}_\gamma(v, \mathcal{T})$ for the γ -height of v in \mathcal{T} . The γ -height of a tree \mathcal{T} , denoted $\text{height}_\gamma(\mathcal{T})$, is defined as $\text{height}_\gamma(r)$, that is the γ -height of its root.

Its name notwithstanding, the definition of γ -height is meant to capture the “bushiness” of a tree. For example, if \mathcal{T} is just a path then its γ -height is equal 0 for each $\gamma \geq 2$. The concept of γ -height generalizes Strahler numbers [23, 24], introduced in hydrology to measure the size of

of their labels. For any node v , when it is its turn to transmit, v transmits any rumor from the set of rumors that have been received so far (including its own rumor) but not yet transmitted. In each round, each rumor that is still not in r will get closer to r , so after n^2 steps all rumors will reach r .

Strong k -selectors. Let $\bar{S} = (S_1, S_2, \dots, S_m)$ be a family of subsets of $\{0, 1, \dots, n-1\}$. \bar{S} is called a *strong k -selector* if, for each k -element set $A \subseteq \{0, 1, \dots, n-1\}$ and each $a \in A$, there is a set S_i such that $S_i \cap A = \{a\}$. As shown in [13, 9], for each k there exists a strong k -selector $\bar{S} = (S_0, S_1, \dots, S_{m-1})$ with $m = O(k^2 \log n)$. We will make extensive use of strong k -selectors in our algorithm. At a certain time in the computation our protocols will “run” \bar{S} , for an appropriate choice of k , by which we mean that it will execute a sequence of m consecutive steps, such that in the j th step the nodes from S_j will transmit, while those not in S_j will stay quiet. This will guarantee that, for any node v with at most $k-1$ siblings, there will be at least one step in the execution of \bar{S} where v will transmit but none of its siblings will. Therefore at least one of v ’s transmissions will be successful.

3 An $O(n\sqrt{\log n})$ -Time Protocol

We first give a gathering protocol SIMPLEGATHER for trees with running time $O(n\sqrt{\log n})$. Our faster protocol will be presented in the next section. We fix three parameters:

$$K = 2^{\lfloor \sqrt{\log n} \rfloor}, D = \lceil \log_K n \rceil = O(\sqrt{\log n}), D' = \lceil \log K^3 \rceil = O(\sqrt{\log n}).$$

We also fix a strong K -selector $\bar{S} = (S_0, S_1, \dots, S_{m-1})$, where $m \leq CK^2 \log n$, for some integer constant C .

By Lemma 1, we have that $\text{height}_K(\mathcal{T}) \leq D$. We call a node v of \mathcal{T} *light* if $|\mathcal{T}_v| \leq n/K^3$; otherwise we say that v is *heavy*. Let \mathcal{T}' be the subtree of \mathcal{T} induced by the heavy nodes. By the definition of heavy nodes, \mathcal{T}' has at most K^3 leaves, so $\text{height}_2(\mathcal{T}') \leq D'$. Also, obviously, $r \in \mathcal{T}'$.

To streamline the description of our algorithm we will allow each node to receive and transmit messages at the same time. We will also assume a preprocessing step allowing each v to know both the size of its subtree \mathcal{T}_v (in particular, whether it is in \mathcal{T}' or not), its K -height, and, if it is in \mathcal{T}' , its 2-height in the subtree \mathcal{T}' . We later explain both the preprocessing and how to modify the algorithm to remove the receive/transmit assumption.

The algorithm consists of two epochs. Epoch 1 consists of $D+1$ stages, each lasting $O(n)$ steps. In this epoch only light vertices participate. The purpose of this epoch is to gather all rumors from \mathcal{T} in \mathcal{T}' . Epoch 2 has $D'+1$ stages and only heavy vertices participate in the computation during this epoch. The purpose of epoch 2 is to deliver all rumors from \mathcal{T}' to r . We describe the computation in the two epochs separately.

A detailed description of Algorithm SIMPLEGATHER is given in Pseudocode 1. To distinguish between computation steps (which do not consume time) and communication steps, we use command “**at time t** ”. When the algorithm reaches this command it waits until time step t to continue processing. Each message transmission takes one time step. For each node v we maintain a set B_v of rumors received by v , including its own rumor ρ_v .

Epoch 1: light vertices. Let $0 \leq h \leq D$, and let v be a light vertex whose K -height equals h . Then v will be active only during stage h which starts at time $\alpha_h = (C+1)hn$. This stage is divided into two parts.

In the first part of stage h , v will transmit according to the strong K -selector \bar{S} . Specifically, this part has n/K^3 iterations, where each iteration corresponds to a complete execution of \bar{S} . At

any time, some of the rumors in B_v may be marked; the marking on a rumor indicates that the algorithm has already attempted to transmit it using \bar{S} . At the beginning of each iteration, v chooses any rumor $\rho_z \in B_v$ it has not yet marked, then transmits ρ_z in the steps that use sets S_i containing the label of v . This ρ_z is then marked. If the parent u of v has degree at most K , the definition of strong K -selectors guarantees that ρ_z will be received by u , but if u 's degree is larger it may not have received ρ_z . Note that the total number of steps required for this part of stage h is $(n/K^3) \cdot m \leq Cn$, so these steps will be completed before the second part of stage h starts.

In the second part, that starts at time $\alpha_h + Cn$, we simply run a variant of the ROUNDROBIN protocol, but cycling through rumors instead of nodes: in the l -th step of this part, all nodes holding the rumor of the node with label l transmit that rumor (note that due to the tree topology it is impossible for two siblings to both be holding rumor ℓ).

We claim that the following invariant holds for all $h = 0, 1, \dots, D$:

(**I_h**) Let $w \in \mathcal{T}$ and let u be a light child of w with $\text{height}_K(u) \leq h - 1$. Then at time α_h node w has received all rumors from \mathcal{T}_u .

To prove this invariant we proceed by induction on h . If $h = 0$ the invariant (**I₀**) holds vacuously. So suppose that invariant (**I_h**) holds for some value of h . We want to prove that (**I_{h+1}**) is true when stage $h + 1$ starts. We thus need to prove the following claim: if u is a light child of w with $\text{height}_K(u) \leq h$ then at time α_{h+1} all rumors from \mathcal{T}_u will arrive in w .

If $\text{height}_K(u) \leq h - 1$ then the claim holds, immediately from the inductive assumption (**I_h**). So assume that $\text{height}_K(u) = h$. Consider the subtree H rooted at u and containing all descendants of u whose K -height is equal to h . By the inductive assumption, at time α_h any $w' \in H$ has all rumors from the subtrees rooted at its descendants of K -height smaller than h , in addition to its own rumor $\rho_{w'}$. Therefore all rumors from \mathcal{T}_u are already in H and each of them has exactly one copy in H , because all nodes in H were idle before time α_h .

When the algorithm executes the first part of stage h on H , then each v node in H whose parent is also in H will successfully transmit an unmarked rumor during each pass through the strong K -selector – indeed, our definition of H guarantees that v has at most $K - 1$ siblings in H , so by the definition of strong selector it must succeed at least once. We make the following additional claim:

Claim 1. *At all times during stage h , the collection of nodes in H still holding unmarked rumors forms an induced tree of H*

The claim follows from induction: At the beginning of the stage the nodes in H still hold their own original rumor, and it is unmarked since those nodes were idle so far. As the stage progresses, each parent of a transmitting child will receive a new (and therefore not yet marked) rumor during each run through the strong selector, so no holes can ever form.

In particular, node u will receive a new rumor during every run through the strong selector until it has received all rumors from its subtree. Since the tree originally held at most $|\mathcal{T}_u| \leq n/K^3$ rumors originally, u must have received all rumors from its subtree after at most n/K^3 runs through the selector.

Note that, as $\text{height}_K(u) = h$, u will also attempt to transmit its rumors to w during this part, but, since we are not making any assumptions about the degree of w , there is no guarantee that w will receive them. This is where the second part of this stage is needed. Since in the second part each rumor is transmitted without collisions, all rumors from u will reach w before time α_{h+1} , completing the inductive step and the proof that (**I_{h+1}**) holds.

In particular, using Invariant (**I_h**) for $h = D$, we obtain that after epoch 1 each heavy node w will have received rumors from the subtrees rooted at all its light children. Therefore at that time all rumors from \mathcal{T} will be already in \mathcal{T}' , with each rumor having exactly one copy in \mathcal{T}' .

Pseudocode 1 SIMPLEGATHER(v)

```

1:  $K = 2^{\lfloor \sqrt{\log n} \rfloor}$ ,  $D = \lceil \log_K n \rceil$ 
2:  $B_v \leftarrow \{\rho_v\}$  ▷ Initially  $v$  has only  $\rho_v$ 
3: Throughout: all rumors received by  $v$  are automatically added to  $B_v$ 
4: if  $|T_v| \leq n/K^3$  then ▷  $v$  is light (epoch 1)
5:    $h \leftarrow \text{height}_K(v, \mathcal{T})$ ;  $\alpha_h \leftarrow (C+1)nh$  ▷  $v$  participates in stage  $h$ 
6:   for  $i = 0, 1, \dots, n/K^3 - 1$  do ▷ iteration  $i$ 
7:     at time  $\alpha_h + im$ 
8:     if  $B_v$  contains an unmarked rumor then ▷ Part 1: strong  $K$ -selector
9:       choose any unmarked  $\rho_z \in B_v$  and mark it
10:      for  $j = 0, 1, \dots, m-1$  do
11:        at time  $\alpha_h + im + j$ 
12:        if  $\text{label}(v) \in S_j$  then TRANSMIT( $\rho_z$ )
13:      for  $l = 0, 1, \dots, n-1$  do ▷ Part 2: RoundRobin
14:        at time  $\alpha_h + Cn + l$ 
15:         $z \leftarrow \text{node with label}(z) = l$ 
16:        if  $\rho_z \in B_v$  then TRANSMIT( $\rho_z$ )
17: else ▷  $v$  is heavy (epoch 2)
18:    $g \leftarrow \text{height}_2(v, \mathcal{T}')$ ;  $\alpha'_g \leftarrow \alpha_{D+1} + 2ng$  ▷  $v$  participates in stage  $g$ 
19:   for  $i = 0, 1, \dots, n-1$  do ▷ Part 1: all nodes transmit
20:     at time  $\alpha'_g + i$ 
21:     if  $B_v$  contains an unmarked rumor then
22:       choose any unmarked  $\rho_z \in B_v$  and mark it
23:       TRANSMIT( $\rho_z$ )
24:   for  $l = 0, 1, \dots, n-1$  do ▷ Part 2: RoundRobin
25:     at time  $\alpha'_g + 2n + l$ 
26:      $z \leftarrow \text{node with label}(z) = l$ 
27:     if  $\rho_z \in B_v$  then TRANSMIT( $\rho_z$ )

```




Figure 2: Proving Invariant (\mathbf{I}_h) . Dark-shaded subtrees of \mathcal{T}_u consist of light nodes with K -height at most $h - 1$. H consists of the descendants of u with K -height equal h .

Epoch 2: heavy vertices. In this epoch we have at most $D' + 1$ stages, and only heavy nodes in \mathcal{T}' participate in the computation. When the epoch starts, all rumors are already in \mathcal{T}' . In stage $D + 1 + g$ the nodes in \mathcal{T}' whose 2-height is equal g will participate. Similar to the stages of epoch 1, this stage has two parts and the second part executes ROUNDROBIN, as before. The difference is that now, in the first part, instead of using the strong K -selector, each heavy node will transmit at each of the n steps.

We need to show that r will receive all rumors at the end. The argument is similar as for light vertices, but with a twist, since we do not use selectors now; instead we have steps when all nodes transmit. In essence, we show that each stage reduces by at least one the 2-depth of the minimum subtree of \mathcal{T}' that contains all rumors.

Specifically, we show that the following invariant holds for all $g = 0, 1, \dots, D'$:

(J_g) Let $w \in \mathcal{T}'$ and let $u \in \mathcal{T}'$ be a child of w with $\text{height}_2(u, \mathcal{T}') \leq g - 1$. Then at time α'_g node w has received all rumors from \mathcal{T}_u .

We prove invariant **(J_g)** by induction on g . For $g = 0$, **(J₀)** holds vacuously. Assume that **(J_g)** holds for some g . We claim that **(J_{g+1})** holds right after stage g .

Choose any child u of w with $\text{height}_2(u, \mathcal{T}') \leq g$. If $\text{height}_2(u, \mathcal{T}') \leq g - 1$, we are done, by the inductive assumption. So we can assume that $\text{height}_2(u, \mathcal{T}') = g$. Let P be the subtree of \mathcal{T}' rooted at u and consisting of all descendants of u whose 2-height in \mathcal{T}' is equal g . Then P is simply a path. By the inductive assumption, for each $w' \in P$, all rumors from the subtrees of w' rooted at its children of 2-height at most $g - 1$ are in w' . Thus all rumors from \mathcal{T}_u are already in P . All nodes in P participate in stage g , but their children outside P do not transmit. Therefore each transmission from any node $x \in P - \{u\}$ during stage g will be successful. Due to pipelining, all rumors from P will reach u after the first part of stage g . (This conclusion can also be derived from treating this computation on P as an instance of the token collection game in Section 2, with each step of the transmissions being one step of the game.) In the second part, all rumors from u will be successfully sent to w . So after stage g all rumors from \mathcal{T}_u will be in w , completing the proof that **(J_{g+1})** holds.

Removing simplifying assumptions. At the beginning of this section we made some simplifying assumptions. It still remains to explain how to modify our algorithm so that it works even if these

assumptions do not hold. These modification are similar to those described in [5], but we include them here for the sake of completeness.

First, we assumed a preprocessing step whereby each v knows certain parameters of its subtree \mathcal{T}_v , including the size, its K -height, etc. The justification for this lies in the algorithm from [5] for information gathering in trees with aggregation. Such an algorithm can be modified to compute in linear time any function f such that $f(v)$ is uniquely determined by the values of f on the children of v . The modification is that each node u , when it sends its message (which, in the algorithm from [5] contains all rumors from \mathcal{T}_u), it will instead send the value of $f(u)$. A node v , after it receives all values of $f(u)$ from each child u , will then compute $f(v)$ ¹.

We also assumed that each node can receive and transmit messages at the same time. We now need to modify the algorithm so that it receives messages only in the receiving state and transmits only in the transmitting state. For the ROUNDROBIN steps this is trivial: a node v is in the transmitting state only if it is scheduled to transmit, otherwise it is in the receiving state. For other steps, we will explain the modification for light and heavy nodes separately.

Consider the computation of the light nodes during the steps when they transmit according to the strong selector. Instead of the strong K -selector, we can use the strong $(K + 1)$ -selector, which will not affect the asymptotic running time. When a node v is scheduled to transmit, it enters the transmitting state, otherwise it is in the receiving state. In the proof, where we argue that the message from v will reach its parent, instead of applying the selector argument to v and its siblings, we apply it to the set of nodes consisting of v , its siblings, and its parent, arguing that there will be a step when v is the only node transmitting among its siblings and its parent is in the receiving state.

Finally, consider the computation of the heavy nodes, at steps when all of them transmit. We modify the algorithm so that, in any stage g , the iteration (in Line 18) of these steps is preceded by $O(n)$ -time preprocessing. Recall that the nodes whose 2-height in \mathcal{T}' is equal g form disjoint paths. We can run a one round of ROUNDROBIN where each node transmits an arbitrary message. This way, each node will know whether it is the first node on one of these paths or not. If a node x is first on some path, say P , x sends a message along this path, so that each node $y \in P$ can compute its distance from x . Then, in the part where all nodes transmit, we replace each step by two consecutive steps (even and odd), and we use parity to synchronize the computation along these paths: the nodes at even positions are in the receiving state at even steps and in the transmitting state at odd steps, and the nodes at odd positions do the opposite.

Summarizing this section, we have presented Algorithm SIMPLEGATHER that completes information gathering in ad-hoc radio networks with tree topologies in time $O(n\sqrt{\log n})$. In the next section, we will show how to improve this bound to $O(n \log \log n)$.

4 A Protocol with Running Time $O(n \log \log n)$

In this section we consider the same basic model of information gathering in trees as in Section 3, that is, the model does not provide any collision detection and rumor aggregation is not allowed. We show that how to refine our $O(n\sqrt{\log n})$ protocol SIMPLEGATHER to improve the running time to $O(n \log \log n)$. This is the first main result of our paper, as summarized in the theorem below.

Theorem 1. *The problem of information gathering on trees, without rumor aggregation, can be solved in time $O(n \log \log n)$.*

¹It needs to be emphasized here that in our model only communication steps contribute to the running time; all calculations are assumed to be instantaneous.

Our protocol that achieves running time $O(n \log \log n)$ will be called FASTGATHER. This protocol can be thought of as an iterative application of the idea behind Algorithm SIMPLEGATHER from Section 3. We assume that the reader is familiar with Algorithm SIMPLEGATHER and its analysis, and in our presentation we will focus on the high level ideas behind Algorithm FASTGATHER, referring the reader to Section 3 for the implementation of some details.

As before, we use notation \mathcal{T} for the input tree and $n = |\mathcal{T}|$ denotes the number of vertices in \mathcal{T} . We assume that n is sufficiently large, and we will establish some needed lower bounds for n as we work through the proof. We fix some arbitrary integer constant $\beta \geq 2$. For $\ell = 1, 2, \dots$, let $K_\ell = \lceil n^{\beta^{-\ell}} \rceil$. So $K_1 = \lceil n^{1/\beta} \rceil$, the sequence $(K_\ell)_\ell$ is non-increasing, and $\lim_{\ell \rightarrow \infty} K_\ell = 2$. Let L be the largest value of ℓ for which $n^{\beta^{-\ell}} \geq \log n$. (Note that L is well defined for sufficiently large n , since β is fixed). We thus have the following exact and asymptotic bounds:

$$\begin{aligned} L &\leq \log_\beta(\log n / \log \log n) & K_L &\geq \log n \\ L &= \Theta(\log \log n) & K_L &= \Theta(\log n). \end{aligned}$$

For $\ell = 1, 2, \dots, L$, by $\bar{S}^\ell = (S_1^\ell, S_2^\ell, \dots, S_{m_\ell}^\ell)$ we denote a strong K_ℓ -selector of size $m_\ell \leq CK_\ell^2 \log n$, for some integer constant C . As discussed in Section 2, such selectors \bar{S}^ℓ exist.

Let $\mathcal{T}^{(0)} = \mathcal{T}$, and for each $\ell = 1, 2, \dots, L$, let $\mathcal{T}^{(\ell)}$ be the subtree of \mathcal{T} induced by the nodes v with $|\mathcal{T}_v| \geq n/K_\ell^3$. Each tree $\mathcal{T}^{(\ell)}$ is rooted at r , and $\mathcal{T}^{(\ell)} \subseteq \mathcal{T}^{(\ell-1)}$ for $\ell \geq 1$. For $\ell \neq 0$, the definition of $\mathcal{T}^{(\ell)}$ implies also that it has at most K_ℓ^3 leaves, so, by Lemma 1, its $K_{\ell+1}$ -height is at most $\log_{K_{\ell+1}}(K_\ell^3)$. Since $K_\ell \leq 2n^{\beta^{-\ell}}$ and $K_{\ell+1} \geq n^{\beta^{-(\ell+1)}}$, we have

$$\begin{aligned} \log_{K_{\ell+1}}(K_\ell^3) &= 3 \log_{K_{\ell+1}}(K_\ell) \\ &\leq 3 \log_{n^{\beta^{-(\ell+1)}}}(2n^{\beta^{-\ell}}) \\ &\leq 3 \log_{n^{\beta^{-(\ell+1)}}} n^{\beta^{-\ell}} + 3 \log_{n^{\beta^{-(\ell+1)}}} 2 \\ &= 3\beta + \frac{3\beta^{\ell+1}}{\log n} \leq 3\beta + \frac{3\beta^{L+1}}{\log n} \leq 3\beta + 1, \end{aligned}$$

where the last inequality holds as long as $3\beta^{L+1} \leq \log n$, which is true if n is large enough (since by the first of the "exact and asymptotic bounds" above we have $\beta^L \leq \log n / \log \log n$). We thus obtain that the $K_{\ell+1}$ -height of $\mathcal{T}^{(\ell)}$ is at most $D = 3\beta + 1 = O(1)$.

As in the previous section, we will make some simplifying assumptions. First, we will assume that all nodes can receive and transmit messages at the same time. Second, we will also assume that each node v knows the size of its subtree $|\mathcal{T}_v|$ and its K_ℓ -heights, for each $\ell \leq L$. After completing the description of the algorithm we will explain how to modify it so that these assumptions will not be needed.

Algorithm FASTGATHER consists of $L + 1$ epochs, numbered $1, 2, \dots, L + 1$. In each epoch $\ell \leq L$ only the nodes in $\mathcal{T}^{(\ell-1)} - \mathcal{T}^{(\ell)}$ participate. At the beginning of epoch ℓ all rumors will be already collected in $\mathcal{T}^{(\ell-1)}$, and the purpose of epoch ℓ is to move all these rumors into $\mathcal{T}^{(\ell)}$. Each of these L epochs will run in time $O(n)$, so their total running time will be $O(nL) = O(n \log \log n)$. In epoch $L + 1$, only the nodes in $\mathcal{T}^{(L)}$ participate. At the beginning of this epoch all rumors will be already in $\mathcal{T}^{(L)}$, and when this epoch is complete, all rumors will be collected in r . This epoch will take time $O(n \log \log n)$. Thus the total running time will be also $O(n \log \log n)$. We now provide the details.

Epochs $\ell = 1, 2, \dots, L$. In epoch ℓ , only the nodes in $\mathcal{T}^{(\ell-1)} - \mathcal{T}^{(\ell)}$ are active, all other nodes will be quiet. The computation in this epoch is very similar to the computation of light nodes (in

epoch 1) in Algorithm SIMPLEGATHER. Epoch ℓ starts at time $\gamma_\ell = (D+1)(C+1)(\ell-1)n$ and lasts $(D+1)(C+1)n$ steps.

Let $v \in \mathcal{T}^{(\ell-1)} - \mathcal{T}^{(\ell)}$. The computation of v in epoch ℓ consists of $D+1$ identical stages. Each stage $h = 0, 1, \dots, D$ starts at time step $\alpha_{\ell,h} = \gamma_\ell + (C+1)hn$ and lasts $(C+1)n$ steps.

Stage h has two parts. The first part starts at time $\alpha_{\ell,h}$ and lasts time Cn . During this part we execute $\lceil n/K_\ell^3 \rceil$ iterations, each iteration consisting of running the strong $K_{\ell+1}$ -selector \bar{S}^ℓ . The time needed to execute these iterations is at most

$$\begin{aligned} \lceil \frac{n}{K_\ell^3} \rceil \cdot CK_{\ell+1}^2 \log n &\leq Cn \cdot \frac{2K_{\ell+1}^2 \log n}{K_\ell^3} \\ &\leq Cn \cdot \frac{8n^{2\beta^{-(\ell+1)}} \log n}{n^{3\beta^{-\ell}}} \\ &= Cn \cdot \frac{8 \log n}{n^{(3-2/\beta)\beta^{-\ell}}} \\ &\leq Cn \cdot \frac{8 \log n}{n^{(3-2/\beta)\beta^{-L}}} \leq Cn, \end{aligned}$$

where the last inequality holds as long as $n^{3-2/\beta}\beta^{-L} \geq 8 \log n$, which is again true for sufficiently large n (recall that $\beta \geq 2$ is constant, and $L = \Theta(\log \log n)$).

Thus all iterations executing the strong selector will complete before time $\alpha_{\ell,h} + Cn$. Then v stays idle until time $\alpha_{\ell,h} + Cn$, which is when the second part starts. In the second part we run the ROUNDROBIN protocol, which takes n steps. So stage h will complete right before step $\alpha_{\ell,h} + (C+1)n = \alpha_{\ell,h+1}$. Note that the whole epoch will last $(D+1)(C+1)n$ steps, as needed.

We claim that the algorithm preserves the following invariant for $\ell = 1, 2, \dots, L$:

(\mathbf{I}^ℓ) Let $w \in \mathcal{T}$ and let $u \in \mathcal{T} - \mathcal{T}^{(\ell-1)}$ be a child of w . Then w will receive all rumors from \mathcal{T}_u before time γ_ℓ , that is before epoch ℓ starts.

For $\ell = 1$, invariant (\mathbf{I}^1) holds vacuously, because $\mathcal{T}^{(0)} = \mathcal{T}$. In the inductive step, assume that (\mathbf{I}^ℓ) holds for some epoch ℓ . We want to show that ($\mathbf{I}^{\ell+1}$) holds right after epoch ℓ ends. In other words, we will show that if w has a child $u \in \mathcal{T} - \mathcal{T}^{(\ell)}$ then w will receive all rumors from \mathcal{T}_u before time $\gamma_{\ell+1}$.

So let $u \in \mathcal{T} - \mathcal{T}^{(\ell)}$. If $u \notin \mathcal{T}^{(\ell-1)}$ then ($\mathbf{I}^{\ell+1}$) holds for u , directly from the inductive assumption. We can thus assume that $u \in \mathcal{T}^{(\ell-1)} - \mathcal{T}^{(\ell)}$.

The argument now is very similar to that for Algorithm SIMPLEGATHER in Section 3, when we analyzed epoch 1. For each $h = 0, 1, \dots, D$ we prove a refined version of condition ($\mathbf{I}^{\ell+1}$):

($\mathbf{I}_h^{\ell+1}$) Let $w \in \mathcal{T}$ and let $u \in \mathcal{T}^{(\ell-1)} - \mathcal{T}^{(\ell)}$ be child of w with $\text{height}_{K_\ell}(u, \mathcal{T}^{(\ell-1)}) \leq h-1$. Then w will receive all rumors from \mathcal{T}_u before time $\alpha_{\ell,h}$, that is before stage h .

The proof is the same as for Invariant (\mathbf{I}_h) in Section 3, proceeding by induction on h . For each fixed h we consider a subtree \mathcal{H} rooted at u and consisting of all descendants of u in $\mathcal{T}^{(\ell-1)}$ whose K_ℓ -height is at least h . By the inductive assumption, at the beginning of stage h all rumors from \mathcal{T}_u are already in \mathcal{H} . Then, the executions of \bar{S}^ℓ , followed by the execution of ROUNDROBIN, will move all rumors from \mathcal{H} to w . We omit the details here.

By applying condition ($\mathbf{I}_h^{\ell+1}$) with $h = D$, we obtain that after all stages of epoch ℓ are complete, that is at right before time $\gamma_{\ell+1}$, w will receive all rumors from \mathcal{T}_u . Thus ℓ invariant ($\mathbf{I}^{\ell+1}$) will hold.

Epoch $L+1$. Due to the definition of L , we have that $\mathcal{T}^{(L)}$ contains at most $K_L^3 = O(\log^3 n)$ leaves, so its 2-depth is at most $D' = \log(K_L^3) = O(\log \log n)$, by Lemma 1. The computation in this epoch

is similar to epoch 2 from Algorithm SIMPLEGATHER. As before, this epoch consists of $D' + 1$ stages, where each stage $g = 0, 1, \dots, D'$ has two parts. In the first part, we have n steps in which each node transmits. In the second part, also of length n , we run one iteration of ROUNDROBIN.

Let $\alpha'_g = \gamma_L + 2gn$. To prove correctness, we show that the following invariant holds for all stages $g = 0, 1, \dots, D'$:

(**J** _{g}) Let $w \in \mathcal{T}^{(L)}$ and let $u \in \mathcal{T}^{(L)}$ be a child of w with $\text{height}_2(u, \mathcal{T}^{(L)}) \leq g - 1$. Then at time α'_g node w has received all rumors from \mathcal{T}_u .

The proof is identical to the proof of the analogous Invariant (**J** _{g}) in Section 3, so we omit it here. Applying Invariant (**J** _{g}) with $g = D'$, we conclude that after stage D' , the root r will receive all rumors.

As for the running time, we recall that $L = O(\log \log n)$. Each epoch $\ell = 1, 2, \dots, L$ has $D + 1 = O(1)$ stages, where each stage takes time $O(n)$, so the execution of the first L epochs will take time $O(n \log \log n)$. Epoch $L + 1$ has $D' + 1 = O(\log \log n)$ stages, each stage consisting of $O(n)$ steps, so this epoch will complete in time $O(n \log \log n)$. We thus have that the overall running time of our protocol is $O(n \log \log n)$.

It remains to explain that the simplifying assumptions we made at the beginning of this section are not needed. Computing all subtree sizes and all K_ℓ -heights can be done recursively bottom-up, using the linear-time information gathering algorithm from [5] that uses aggregation. This was explained in Section 3. The difference now is that each node has to compute $L + 1 = O(\log \log n)$ values K_ℓ , and, since we limit bookkeeping information in each message up to $O(\log n)$ bits, these values need to be computed separately. Nevertheless, the total pre-computation time will still be $O(n \log \log n)$.

Removing the assumption that nodes can receive and transmit at the same time can be done in the same way as in Section 3. Roughly, in each epoch $\ell = 1, 2, \dots, L$, any node $v \in \mathcal{T}^{(\ell-1)} - \mathcal{T}^{(\ell)}$ uses a strong $(K_\ell + 1)$ -selector (instead of a strong K_ℓ -selector) to determine whether to be in the receiving or transmitting state. In epoch L the computation (in the steps when all nodes transmit) is synchronized by transmitting a control message along induced paths, and then choosing the receiving or transmitting state according to node parity.

Summarizing, we have proved that the running time of Algorithm SIMPLEGATHER is $O(n \log \log n)$, thus completing the proof of Theorem 1.

5 An $O(n)$ -time Protocol with Acknowledgments

In this section we consider a slightly different communication model from that in Sections 3 and 4. We now assume that acknowledgments of successful transmissions are provided to the sender. All the remaining nodes, including the intended recipient, cannot distinguish between collisions and absence of transmissions. The main result of this section, as summarized in the theorem below, is that with this feature it is possible to achieve the optimal time $O(n)$.

Theorem 2. *The problem of information gathering on trees without rumor aggregation can be solved in time $O(n)$ if acknowledgments are provided.*

The overall structure of our $O(n)$ time protocol, called Algorithm LINGATHER, is similar to Algorithm SIMPLEGATHER. It consists of two epochs. The first epoch does not use the acknowledgement feature, and it is in fact identical to Epoch 1 in Algorithm SIMPLEGATHER, except for a

different choice of the parameters. After this epoch, lasting time $O(n)$, all rumors will be collected in the subtree \mathcal{T}' consisting of the heavy nodes.

In the second epoch, only the heavy nodes in \mathcal{T}' will participate in the computation, and the objective of this epoch is to move all rumors already collected in \mathcal{T}' to the root r . The key obstacle to be overcome in this epoch is congestion stemming from the fact that, although \mathcal{T}' may be small, its nodes have many rumors to transmit. This congestion means that simply repeatedly applying k -selectors is no longer enough. For example, if the root has ℓ children, each with n/ℓ rumors, then repeating an ℓ -selector n/ℓ times would get all the rumors to the root. However, we know from the selector size bounds in [8] that this would take total time $\Omega(n\ell \log n / \log \ell)$, which is far too slow. While in this particular scenario ROUNDROBIN will collect all rumors in the root in time $O(n)$, this example can be enhanced to show that simple combinations of k -selectors and ROUNDROBIN do not seem to be sufficient to gather all rumors from \mathcal{T}' in the root in linear time.

To overcome this obstacle, we introduce two novel tools that will play a critical role in our algorithm. The first tool is a so-called *amortizing selector family*. Since a parent, say with ℓ children, receives at most one rumor per round, it clearly cannot simultaneously be receiving rumors at an average rate greater than $\frac{1}{\ell}$ from each child individually. With the amortizing family, we will be able to achieve this bound within a constant fraction over long time intervals, so long as each child knows (approximately) how many siblings it is competing with.

Of course, such a family will not be useful unless a node can obtain an accurate estimate of its parent's degree, which will be the focus of our second tool, *k-distinguishers*. Using a k -distinguisher a node can determine whether its number of active siblings ℓ is at least k or at most $2k$. While this information is not sufficient to determine the exact relation between ℓ and k , we show how to combine different k -distinguishers to obtain another structure, called a *cardinality estimator*, that will determine the value of ℓ within a factor of 4. Using this estimate, and the earlier described amortizing selector, a node can quickly transmit its rumors to its parent. This will allow us to gather all rumors from \mathcal{T}' in the root in time $O(n)$.

This section is divided into three parts. In Sections 5.1 and 5.2 we give precise definitions and constructions of our combinatorial structures. Our $O(n)$ -time protocol LINGATHER is described and analyzed in Section 5.3.

5.1 Construction of Amortizing Selectors

We now define the concept of an *amortizing selector family* \bar{S} . Similarly to a strong selector, this amortizing family will be a collection of subsets of the underlying label set $[n]$, though now it will be doubly indexed. Specifically, $\bar{S} = \{S_{ij}\}$, where $1 \leq i \leq s$ and each $j \in \{1, 2, 4, 8, \dots, k\}$, for some parameters s and k . We say that \bar{S} *succeeds at cumulative rate* q if the following statement is true:

For each $j \in \{1, 2, 4, \dots, \frac{k}{2}\}$, each subset $A \subseteq \{1, \dots, n\}$ satisfying $j/2 \leq |A| \leq 2j$, and each element $v \in A$ there are at least $\frac{q}{|A|}s$ distinct i for which

$$v \in S_{ij} \text{ and } A \cap (S_{i(j/2)} \cup S_{ij} \cup S_{i(2j)}) = \{v\}.$$

In the case $j = 1$ the set $S_{i(j/2)}$ is defined to be empty. Here s can be thought of as the total running time of the selector, j as a node's estimate of its parent's degree, and k as some bound on the maximum degree handled by the selector. A node fires at time step i if and only if its index is contained in the set S_{ij} . What the above statement is then saying that for any subset A of siblings, if $|A|$ is at most $k/2$ and each child estimates $|A|$ within a factor of 2 then each child will transmit at rate at least $\frac{q}{|A|}$.

Theorem 3. *There are fixed constants $c, C > 0$ such that the following is true: For any k and n and any $s \geq Ck^2 \log n$, there is an amortizing selector with parameters n, k, s succeeding with cumulative rate c .*

Let k, n, s be given such that $s \geq Ck^2 \log n$, where k is a constant to be determined later. We form our selector probabilistically: For each v, i , and j , we independently include v in S_{ij} with probability 2^{-j} .

Observe that by monotonicity it suffices to check the selector property for the case $|A| = 2j$: If we replace A with a larger set A' containing A and satisfying $|A'| \leq 4|A|$, then for any $v \in A$ and any i satisfying

$$A' \cap (S_{i(j/2)} \cup S_{ij} \cup S_{i(2j)}) = \{v\},$$

we also have

$$A \cap (S_{i(j/2)} \cup S_{ij} \cup S_{i(2j)}) = \{v\}.$$

So if there are at least $\frac{cs}{|A'|}$ distinct i satisfying the first equality, there are at least $\frac{(c/4)s}{|A|}$ satisfying the second equality.

Now fix $j \in \{1, 2, 4, \dots, k\}$, a set $A \subseteq [n]$ with $|A| = 2j$ and some $v \in A$, and let the random variables X and Y be defined by

$$X = |\{i : v \in S_{ij} \text{ and } A \cap (S_{i(j/2)} \cup S_{ij} \cup S_{i(2j)}) = \{v\}\}|$$

The expected value of X is

$$\mu_X = s \left(\frac{1}{j}\right) \left(1 - \frac{2}{j}\right)^{2j-1} \left(1 - \frac{1}{j}\right)^{2j-1} \left(1 - \frac{1}{2j}\right)^{2j-1}.$$

Utilizing the bound

$$\left(1 - \frac{2}{j}\right)^{2j-1} \left(1 - \frac{1}{j}\right)^{2j-1} \left(1 - \frac{1}{2j}\right)^{2j-1} \geq \frac{1}{e^7},$$

we have

$$\mu_X \geq e^{-7} \frac{s}{j}$$

Now let $c = \frac{1}{4e^7}$. Applying the Chernoff bound, we get

$$\begin{aligned} \Pr[X \leq c \frac{s}{|A|}] &\leq \Pr[X \leq \frac{1}{2} \mu_X] \leq e^{-\mu_X/8} \\ &\leq e^{-s/8e^7 j}. \end{aligned}$$

We now use the union bound over all choices of j, v , and S . We have at most n choices of v , at most $\log n$ choices for j , and at most $\binom{n}{2j-1} \leq n^{2k-1}$ choices of S given j and v . Thus the probability that our family is not an amortizing selector is at most

$$n \log n \cdot n^{2k-1} \cdot \left(e^{-s/8e^7 j}\right) \leq 2n^{2k} \log n \cdot e^{-Ck \log n / 8e^7}$$

which is smaller than 1 for sufficiently large C . This implies the existence of the amortized selector family.

5.2 Construction of k -Distinguishers and Cardinality Estimators

In this section, we define k -distinguishers and show how to construct them. Let $\bar{S} = (S_1, S_2, \dots, S_m)$, where $S_j \subseteq [n]$ for each j . For $A \subseteq [n]$ and $a \in A$, define $Hits_{a,A}(\bar{S}) = \{j : S_j \cap A = \{a\}\}$, that is $Hits_{a,A}(\bar{S})$ is the collection of indices j for which S_j intersects A exactly on a . Note that, using this terminology, \bar{S} is a strong k -selector if and only if $Hits_{a,A}(\bar{S}) \neq \emptyset$ for all sets $A \subseteq [n]$ of cardinality at most k and all $a \in A$.

We say that \bar{S} is a k -*distinguisher* if there is a *threshold* value ξ (which may depend on k) such that, for any $A \subseteq [n]$ and $a \in A$, the following conditions hold:

$$\text{if } |A| \leq k \text{ then } |Hits_{a,A}(\bar{S})| > \xi, \text{ and if } |A| \geq 2k \text{ then } |Hits_{a,A}(\bar{S})| < \xi.$$

We make no assumptions on what happens for $|A| \in \{k+1, k+2, \dots, 2k-1\}$.

The idea is this: consider a fixed a , and imagine that we have some set A that contains a , but its other elements are not known. Suppose that we also have an access to a *hit oracle* that for any set S will tell us whether $S \cap A = \{a\}$ or not. With this oracle, we can then use a k -distinguisher \bar{S} to extract some information about the cardinality of A by calculating the cardinality of $Hits_{a,A}(\bar{S})$. If $|Hits_{a,A}(\bar{S})| \leq \xi$ then we know that $|A| > k$, and if $|Hits_{a,A}(\bar{S})| \geq \xi$ then we know that $|A| < 2k$. What we will soon show, again by a probabilistic construction, is that not-too-large k -distinguishers exist:

Theorem 4. *For any $n \geq 2$ and $1 \leq k \leq n/2$ there exists a k -distinguisher of size $m = O(k^2 \log n)$.*

In our framework, the acknowledgement of a message received from a parent corresponds exactly to such a hit oracle. So if all nodes fire according to such a k -distinguisher, each node can determine in time $O(k^2 \log n)$ either that its parent has at least k children or that it has at most $2k$ children.

Now let λ be a fixed parameter between 0 and 1. For each $i = 0, 1, \dots, \lceil \lambda \log n \rceil$, let \bar{S}^i be a 2^i -distinguisher of size $O(2^{2i} \log n)$ and with threshold value ξ_i . We can then concatenate these k -distinguishers to obtain a sequence \tilde{S} of size $\sum_{i=0}^{\lceil \lambda \log n \rceil} O(2^{2i} \log n) = O(n^{2\lambda} \log n)$.

We will refer to \tilde{S} as a *cardinality estimator*, because applying our hit oracle to \tilde{S} we can estimate a cardinality of an unknown set within a factor of 4, making $O(n^{2\lambda} \log n)$ hit queries. More specifically, consider again a scenario where we have a fixed a and some unknown set A containing a , where $|A| \leq n^\lambda$. Using the hit oracle, compute the values $h_i = |Hits_{a,A}(\bar{S}^i)|$, for all i . If i_0 is the smallest i for which $h_i > \xi_i$, then by definition of our distinguisher we must have $2^{i_0-1} < |A| < 2^{i_0}$. In our gathering framework, this corresponds to each node in the tree being able to determine in time $O(n^{2\lambda} \log n)$ a value of j (specifically, $i_0 - 1$) such that the number of children of its parent is between 2^j and 2^{j+2} , which is exactly what we need to be able to run the amortizing selector.

It remains to show the existence of k -distinguishers, i.e. to prove Theorem 4. Let $m = Ck^2 \log n$, where C is some sufficiently large constant whose value we will determine later. We choose the collection of random sets $\bar{S} = (S_1, S_2, \dots, S_m)$, by letting each S_j be formed by independently including each $x \in [n]$ in S_j with probability $1/2k$. Thus, for any set A and $a \in A$, the probability that $S_j \cap A = \{a\}$ is $(1/2k)(1 - 1/2k)^{|A|-1}$, and the expected value of $|Hits_{a,A}(\bar{S})|$ is

$$\mathbf{E}[|Hits_{a,A}(\bar{S})|] = m \cdot \frac{1}{2k} \left(1 - \frac{1}{2k}\right)^{|A|-1}. \quad (1)$$

Recall that to be a k -distinguisher our set needs to satisfy (for suitable ξ) the following two properties:

- (d1) if $|A| \leq k$ then $|Hits_{a,A}(\bar{S})| > \xi$
- (d2) if $|A| \geq 2k$ then $|Hits_{a,A}(\bar{S})| < \xi$.

We claim that, for a suitable value of ξ , the probability that there exists a set $A \subseteq [n]$ and some $a \in A$ for which \bar{S} does not satisfy both conditions is smaller than 1 (and in fact tends to 0). This will be sufficient to show an existence of a k -distinguisher with threshold value ξ .

Observe that in order to be a k -distinguisher it is sufficient that \bar{S} satisfies (d1) for sets A with $|A| = k$ and satisfies (d2) for sets A with $|A| = 2k$. This is true because the value of $|Hits_{a,A}(\bar{S})|$ is monotone with respect to the inclusion: if $a \in A \subseteq B$ then $Hits_{a,A}(\bar{S}) \supseteq Hits_{a,B}(\bar{S})$.

Now consider some fixed $a \in [n]$ and two sets $A_1, A_2 \subseteq [n]$ such that $|A_1| = k$, $|A_2| = 2k$ and $a \in A_1 \cap A_2$. For $i = 1, 2$, we consider two corresponding random variables $X_i = |Hits_{a,A_i}(\bar{S})|$ and their expected values $\mu_i = \text{Exp}[X_i]$. For any integer $k \geq 1$ we have

$$\begin{aligned} \frac{1}{e^{1/2}} &\leq \left(1 - \frac{1}{2k}\right)^{k-1} \\ \frac{1}{e} &\leq \left(1 - \frac{1}{2k}\right)^{2k-1} \leq \frac{1}{2}. \end{aligned}$$

From (1), substituting $m = Ck^2 \log n$, this gives us the corresponding estimates for μ_1 and μ_2 :

$$\begin{aligned} \frac{1}{2e^{1/2}} Ck \log n &\leq \mu_1 \\ \frac{1}{2e} Ck \log n &\leq \mu_2 \leq \frac{1}{4} Ck \log n \end{aligned}$$

Since $e^{-1/2} > \frac{1}{2}$, we can choose an $\epsilon \in (0, 1)$ and ξ for which

$$(1 + \epsilon)\mu_2 < \xi < (1 - \epsilon)\mu_1.$$

Thus the probability that \bar{S} violates (d1) for $A = A_1$ is

$$\begin{aligned} \Pr[X_1 \leq \xi] &\leq \Pr[X_1 \leq (1 - \epsilon)\mu_1] \leq e^{-\epsilon^2 \mu_1 / 2} \\ &\leq e^{-\epsilon^2 e^{-1/2} Ck \log n / 4}, \end{aligned}$$

where in the second inequality we use the Chernoff bound for deviations below the mean. Similarly, using the Chernoff bound for deviations above the mean, we can bound the probability of \bar{S} violating (d2) for $A = A_2$ as follows:

$$\begin{aligned} \Pr[X_2 \geq \xi] &\leq \Pr[X_2 \geq (1 + \epsilon)\mu_2] \leq e^{-\epsilon^2 \mu_2 / 3} \\ &\leq e^{-\epsilon^2 e^{-1} Ck \log n / 6}. \end{aligned}$$

To finish off the proof, we apply the union bound. We have at most n choices for a , at most $\binom{n}{k-1} \leq n^{k-1} \leq n^{2k-1}$ choices of A_1 , and at most $\binom{n}{2k-1} \leq n^{2k-1}$ choices of A_2 . Note also that $e^{-1/2}/4 > e^{-1}/6$. Putting it all together, the probability that \bar{S} is not a k -distinguisher is at most

$$\begin{aligned} n \cdot \binom{n}{k-1} \cdot \Pr[X_1 \leq \xi] + n \cdot \binom{n}{2k-1} \cdot \Pr[X_2 \geq \xi] &\leq n^{2k} \cdot (\Pr[X_1 \leq \xi] + \Pr[X_2 \geq \xi]) \\ &\leq 2n^{2k} \cdot e^{-\epsilon^2 e^{-1} Ck \log n / 6} \\ &= 2n^{k(2 - \epsilon^2 e^{-1} C/6)} < 1, \end{aligned}$$

for C large enough.

5.3 Linear-time Protocol

As before, \mathcal{T} is the input tree with n nodes. We will recycle the notions of light and heavy nodes from Section 3, although now we will use slightly different parameters. Let $\delta > 0$ be a small constant, and let $K = \lceil n^\delta \rceil$. We say that $v \in \mathcal{T}$ is *light* if $|T_v| \leq n/K^3$ and we call v *heavy* otherwise. By \mathcal{T}' we denote the subtree of \mathcal{T} induced by the heavy nodes.

Algorithm LINGATHER. Our algorithm will consist of two epochs. The first epoch is essentially identical to Epoch 1 in Algorithm SIMPLEGATHER, except for a different choice of the parameters. The objective of this epoch is to collect all rumors in \mathcal{T}' in time $O(n)$. In the second epoch, only the heavy nodes in \mathcal{T}' will participate in the computation, and the objective of this epoch is to gather all rumors from \mathcal{T}' in the root r . This epoch is quite different from our earlier algorithms and it will use the combinatorial structures obtained in the previous sections to move all rumors from \mathcal{T}' to r in time $O(n)$.

Epoch 1: In this epoch only light nodes will participate, and the objective of Epoch 1 is to move all rumors into \mathcal{T}' . In this epoch we will not be taking advantage of the acknowledgement mechanism. As mentioned earlier, except for different choices of parameters, this epoch is essentially identical to Epoch 1 of Algorithm SIMPLEGATHER, so we only give a very brief overview here. We use a strong K -selector \tilde{S} of size $m \leq CK^2 \log n$.

Let $D = \lceil \log_K n \rceil \leq 1/\delta = O(1)$. By Lemma 1, the K -depth of \mathcal{T} is at most D . Epoch 1 consists of $D + 1$ stages, where in each stage $h = 0, 1, \dots, D$, nodes of K -depth h participate. Stage h consists of n/K^3 executions of \tilde{S} , followed by an execution of ROUNDROBIN, taking total time $n/K^3 \cdot m + n = O(n)$. So the entire epoch takes time $(D + 1) \cdot O(n) = O(n)$ as well. The proof of correctness (namely that after this epoch all rumors are in \mathcal{T}') is identical as for Algorithm SIMPLEGATHER.

Epoch 2: When this epoch starts, all rumors are already gathered in \mathcal{T}' , and the objective is to push them further to the root.

For the second epoch, we restrict our attention to the tree \mathcal{T}' of heavy nodes. As before, no parent in this tree can have more than $K^3 = n^{3\delta}$ children, since each child is itself the ancestor of a subtree of size n/K^3 . We will further assume the existence of a fixed amortizing selector family with parameters $k = 2K^3$ and $s = K^8$, as well as a fixed cardinality estimator with parameter $\lambda = 3\delta$ running in time $D_1 = O(n^{6\delta} \log n) = O(K^6 \log n)$.

Our protocol will be divided into stages, each consisting of $2(D_1 + K^8)$ steps. A node will be *active* in a given stage if at the beginning of the stage it has already received all of its rumors, but still has at least one rumor left to transmit (it is possible for a node to never become active, if it receives its last rumor and then finishes transmitting before the beginning of the next stage).

During each odd-numbered time step of a stage, all nodes (active or not) holding at least one rumor they have not yet successfully passed on transmit such a rumor. The even-numbered time steps are themselves divided into two parts. In the first D_1 even steps, all active nodes participate in the aforementioned cardinality estimator. At the conclusion of the estimator, each node knows a j such that their parent has between 2^j and 2^{j+2} active children. Note that active siblings do not necessarily have the same estimate for their family size. For the remainder of the even steps, each active node fires using the corresponding 2^{j+1} -selector from the amortizing family.

The stages repeat until all rumors have reached the root. Our key claim, stated below, is that the rumors aggregate at least at a steady rate over time – each node with subtree size m in the original tree \mathcal{T} will have received all m rumors within $O(m)$ steps of the start of the epoch.

(\mathbf{I}^ℓ) For any heavy node v such that v has subtree size m in \mathcal{T} , and any $0 \leq j \leq m$, that node

has received at least j rumors within time $C(2m + j)$ of the beginning of Epoch 2, where C is some sufficiently large absolute constant.

In particular, the root has received all of the rumors by time $3Cn$.

We will show this invariant holds by induction on the node's height within \mathcal{T}' . If the node is a leaf the statement follows from our analysis of Epoch 1 (the node has received *all* rumors from its subtree by the beginning of epoch 2).

Now assume that a node u with subtree size $m + 1$ has k children within \mathcal{T}' , and that those children have subtree sizes $a_1 \geq a_2 \geq \dots \geq a_k \geq K^3$. Node u may also have received some number a_0 of messages from non-heavy children (these children, if any, will have already completed their transmissions during the previous epoch).

Let v be a child having subtree size a_1 (chosen arbitrarily, if there are two children with maximal subtree size). Let t_2 be defined by

$$t_2 = \begin{cases} 3Ca_2 + \frac{3}{c}(a_2 + \dots + a_k) + K^{12} & \text{if } k > 1 \\ 0 & \text{if } k = 1 \end{cases}$$

We make the following additional claims.

Claim 2: By time t_2 , all children except possibly v will have completed all of their transmissions.

Proof. By inductive hypothesis, all children except v will have received all of their rumors by time $3Ca_2$. During each stage from that point until all non- v nodes complete, one of the following things must be true.

- Some active node transmits the final rumor it is holding, and stops transmitting. This can happen during at most K^3 stages, since there are at most K^3 children and each child only becomes active once it already has received all rumors from its subtree.
- All active nodes have rumors to transmit throughout the entire stage. If there were j active nodes total during the stage, then by the definition of our amortizing selector family, the parent received at least $c\frac{2K^8}{j}$ rumors from each child during the stage. In particular, it must have received at least

$$c(j-1)\frac{2K^8}{j} \geq cK^8$$

new rumors from children other than v .

Combining the two types of stages, the non- v children will have all finished in at most

$$K^3 + \frac{1}{cK^8}(a_2 + \dots + a_k)$$

complete stages after time $3Ca_2$. Since each stage takes time $2(D_1 + K^8) = (2 + o(1))K^8$, the bound follows. \square

Claim 3: Let k, m , and v be as above. By time $2Cm$, all children except possibly v have completed their transmissions.

Proof. This is trivial for $k = 1$. For larger k follows from the previous claim, together with the estimate that (for sufficiently large C)

$$2Cm \geq 2C(a_1 + \dots + a_k) \geq 4Ca_2 + 2C(a_3 + \dots + a_k) \geq \frac{4}{3}(t_2 - K^{12}) \geq \frac{5}{4}t_2$$

Here the middle inequality holds for any $C > 2/c$, while the latter inequality holds since $t_2 \geq a_2 \geq n/K^3 \gg K^{12}$. \square

By the above claim, node v is the only node that could still be transmitting at time $2Cm$. In particular, if it has a rumor during an odd numbered time step after this point, it successfully transmits. By assumption, v will have received at least j rumors by time $C(2m + j)$ for each j . This implies it will successfully broadcast j rumors by time $C(2m + j) + 2$ for each $0 \leq j \leq a_1$.

By time $C[2(m + 1) + j]$, the parent has either received all rumors from all of its children (if $j > a_1$), or at least j rumors from a_1 alone (if $j \leq a_1$). Either way, it has received at least j rumors total, and the induction is complete.

References

- [1] Noga Alon, Amotz Bar-Noy, Nathan Linial, and David Peleg. A lower bound for radio broadcast. *J. Comput. Syst. Sci.*, 43(2):290–298, 1991.
- [2] Danilo Bruschi and Massimiliano Del Pinto. Lower bounds for the broadcast problem in mobile radio networks. *Distrib. Comput.*, 10(3):129–135, April 1997.
- [3] Bogdan S. Chlebus, Leszek Gasieniec, Alan Gibbons, Andrzej Pelc, and Wojciech Rytter. Deterministic broadcasting in ad hoc radio networks. *Distributed Computing*, 15(1):27–38, 2002.
- [4] Malin Christersson, Leszek Gasieniec, and Andrzej Lingas. Gossiping with bounded size messages in ad hoc radio networks. In *Automata, Languages and Programming, 29th International Colloquium, ICALP’02*, pages 377–389, 2002.
- [5] Marek Chrobak, Kevin Costello, Leszek Gasieniec, and Darek R. Kowalski. Information gathering in ad-hoc radio networks with tree topology. In *Proc. 8th International Conference on Combinatorial Optimization and Applications (COCOA’14)*, pages 129–145, 2014.
- [6] Marek Chrobak, Leszek Gasieniec, and Wojciech Rytter. Fast broadcasting and gossiping in radio networks. *J. Algorithms*, 43(2):177–189, 2002.
- [7] Marek Chrobak, Leszek Gasieniec, and Wojciech Rytter. A randomized algorithm for gossiping in radio networks. *Networks*, 43(2):119–124, 2004.
- [8] Andrea E. F. Clementi, Angelo Monti, and Riccardo Silvestri. Selective families, superimposed codes, and broadcasting on unknown radio networks. In *Prof. 12th Annual Symposium on Discrete Algorithms (SODA’01)*, pages 709–718, 2001.
- [9] Andrea E. F. Clementi, Angelo Monti, and Riccardo Silvestri. Distributed broadcast in radio networks of unknown topology. *Theor. Comput. Sci.*, 302(1-3):337–364, 2003.
- [10] Artur Czumaj and Peter Davies. Almost optimal deterministic broadcast in radio networks. *CoRR*, abs/1506.00853, 2015.
- [11] Artur Czumaj and Wojciech Rytter. Broadcasting algorithms in radio networks with unknown topology. *Journal of Algorithms*, 60(2):115 – 143, 2006.
- [12] G. De Marco and D.R. Kowalski. Contention resolution in a non-synchronized multiple access channel. In *2013 IEEE 27th International Symposium on Parallel Distributed Processing (IPDPS)*, pages 525–533, 2013.

- [13] P. Erdős, P. Frankl, and Z. Füredi. Families of finite sets in which no set is covered by the union of r others. *Israel Journal of Mathematics*, 51(1–2):79–89, 1985.
- [14] Antonio Fernández Anta, Miguel A. Mosteiro, and Jorge Ramón Muñoz. Unbounded contention resolution in multiple-access channels. *Algorithmica*, 67(3):295–314, 2013.
- [15] Leszek Gasieniec. On efficient gossiping in radio networks. In *16th Int. Colloquium on Structural Information and Communication Complexity (SIROCCO'09)*, pages 2–14, 2009.
- [16] Leszek Gasieniec, Tomasz Radzik, and Qin Xin. Faster deterministic gossiping in directed ad hoc radio networks. In *Proc. Scandinavian Workshop on Algorithm Theory (SWAT'04)*, pages 397–407, 2004.
- [17] Dariusz R. Kowalski. On selection problem in radio networks. In *Proceedings of the Twenty-fourth Annual ACM Symposium on Principles of Distributed Computing*, PODC '05, pages 158–166, 2005.
- [18] Dariusz R. Kowalski and Andrzej Pelc. Faster deterministic broadcasting in ad hoc radio networks. *SIAM J. Discrete Math.*, 18(2):332–346, 2004.
- [19] Eyal Kushilevitz and Yishay Mansour. An $\Omega(D \log(N/D))$ lower bound for broadcast in radio networks. *SIAM J. Comput.*, 27(3):702–712, 1998.
- [20] Ding Liu and Manoj Prabhakaran. On randomized broadcasting and gossiping in radio networks. In *8th Annual Int. Conference on Computing and Combinatorics (COCOON'02)*, pages 340–349, 2002.
- [21] Gianluca De Marco. Distributed broadcast in unknown radio networks. In *Proc. 19th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA'08)*, pages 208–217, 2008.
- [22] Gianluca De Marco and Dariusz R. Kowalski. Fast nonadaptive deterministic algorithm for conflict resolution in a dynamic multiple-access channel. *SIAM Journal on Computing*, 44(3):868–888, 2015.
- [23] A.N. Strahler. Hypsometric (area-altitude) analysis of erosional topology. *Bull. Geol. Soc. Amer.*, 63:117–1142, 1952.
- [24] X.G. Viennot. A Strahler bijection between Dyck paths and planar trees. *Discrete Mathematics*, 246:317–329, 2003.
- [25] Ying Xu. An $O(n^{1.5})$ deterministic gossiping algorithm for radio networks. *Algorithmica*, 36(1):93–96, 2003.